

MTR 00B0000063

MITRE TECHNICAL REPORT

JPEG 2000 and WSQ Image Compression Interoperability

February 2001

Margaret A. Lepley

Sponsor: DOJ/FBI
Dept. No.: G034

Contract No.: DAAB07-00-C-C201
Project No.: 0700E02X

Approved for public release; distribution unlimited.

©2001 The MITRE Corporation

MITRE
Center for Integrated Intelligence Systems
Bedford, Massachusetts

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE FEB 2001		2. REPORT TYPE		3. DATES COVERED 00-02-2001 to 00-02-2001	
4. TITLE AND SUBTITLE JPEG 2000 and WSQ Image Compression Interoperability				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) MITRE Corporation, 202 Burlington Road, Bedford, MA, 01730-1420				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 52	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

MITRE Department Approval:

/signed/

Joseph L. Howard

MITRE Project Approval:

/signed/

Norman B. Nill

Abstract

This report explores the degree of compatibility between the wavelet-based WSQ fingerprint compression standard and JPEG 2000 compression standard, with a view towards identifying coexistence or potential migration paths. Theoretical comparison of the two standards led to the introduction of three new elements into JPEG 2000. A prototype WSQ-to-JPEG 2000 transcoder, developed exploiting these new elements was used to test the viability of transcoding, examine errors introduced, and measure changes in file size. Results of fingerprint compression with JPEG 2000 Part 1 alone and successive recompression with WSQ and JPEG 2000 Part 1 are also presented.

KEYWORDS: WSQ, JPEG2000, wavelet, image compression, fingerprint, FBI

Executive Summary

This report explores the degree of compatibility between the wavelet-based WSQ fingerprint compression standard and the JPEG 2000 compression standard. WSQ is the FBI-specified fingerprint compression standard that is used throughout the United States criminal justice system and internationally. JPEG 2000 is a new ISO international standard for general wavelet-based image compression that is expected to be widely available in the future.

This study was funded by the FBI as part of MITRE's Image Standards Development support to the FBI, with a view to anticipating the impact of future interactions between JPEG 2000 and WSQ. Two main segments of study occurred. In the first segment, a comparison was made between the two compression techniques, identifying where they are the same and where different. In the second segment, a testbed for converting WSQ files into JPEG 2000 files was created, tested, and analyzed.

Preliminary tests performed on a small set of fingerprint images gave an indication of what could be expected from JPEG 2000 Part 1 applied to fingerprints, and interactions when both compressions are applied sequentially to the same fingerprint (i.e., in different stages of the processing/dissemination/storage chain). Visual inspection showed that JPEG 2000 Part 1 tended to be somewhat blockier/blurrier than WSQ when applied to fingerprints, although the identification and matching capability of the images did not seem to have changed appreciably.

Theoretical comparison showed that JPEG 2000 Part 1 is missing three elements present in WSQ. These elements, if present, would allow compressed files to be easily 'transcoded' from one format to the other with minimal loss in image content. Transcoding is an operation that partially decodes an image in one format and recodes it into another format, while avoiding recomputation of intermediate data common to both algorithms. Due to MITRE's efforts ensuring suitable coding, testing, and discussion within the ISO JPEG committee, these three elements were formally accepted as additions to JPEG 2000 Part 2.

Once JPEG 2000 Part 2 contained the necessary elements, a testbed was created to allow transcoding WSQ files into JPEG 2000 files. Although the transcoding capability was demonstrated, it was discovered that the process inserts a certain amount of error in the resultant images when converting from WSQ to JPEG 2000. This error was measured and, though very small, is outside the FBI's decoder certification specification. Visual inspection of the resultant images showed, however, that the visible artifacts seen with JPEG 2000 Part 1 had been entirely removed and differences from the expected WSQ output were generally not visible.

As a sidelight to the transcoding studies, it was noticed that the JPEG 2000 compressed files were typically at least 10 percent smaller than corresponding WSQ files.

From the studies, we conclude that it will be feasible for WSQ files to be transcoded to JPEG 2000 Part 2 with no visual loss. However, since there are some minor alterations in a small percentage of the pixel values, a separate study would be needed to investigate any impact of the small transcoding difference for images input to an automated fingerprint identification system (AFIS).

Converting WSQ files to JPEG 2000 Part 1 may be a more easily accessible option, due to the more frequent availability of Part 1 encoders and decoders. Such a conversion will cause more degradation in image quality than the Part 2 transcoding, but the quality may still be high enough for some applications.

WSQ and JPEG 2000 are similar enough that questions may emerge about migration of the FBI standard. JPEG 2000 Part 1 by itself has demonstrably lower visual quality, so it would not be a good alternative. There is some indication from the results shown here that JPEG 2000 Part 2 by itself would be able to achieve similar image quality to WSQ at a slightly smaller file size. However, a small improvement in file size must be weighed against other disadvantages of changing an accepted standard that is already in wide use.

Acknowledgments

This study was funded by the FBI as part of their Image Standards Development project with MITRE. Also acknowledged is MITRE's internal support for the JPEG 2000 standardization effort, which enabled access to JPEG 2000 test code and provided the means for influencing the standard's content in a way that promoted interoperability with the FBI's compression standard.

Table of Contents

Section	Page
1. Introduction	1
1.1 Background on JPEG 2000	1
1.2 Background on WSQ and FBI Involvement	2
1.3 Technical Background	2
1.4 Project Summary	2
2. Interoperability Analysis	3
2.1 WSQ Gray-Scale Fingerprint Image Compression Standard	3
2.2 JPEG 2000	4
2.2.1 Part 1	4
2.2.2 Part 2	5
2.3 Evaluation	5
2.4 Compression Performance	7
2.4.1 Image Metrics	9
2.4.2 Visual Performance	12
2.5 Comments on VM7 vs. WSQ	12
2.6 Conclusions	13
3. Transcoding	15
3.1 Notation	15
3.2 Conversion	15
3.2.1 Subband Ordering	18
3.3 Parameter Format/Precision	19
3.4 Implementation Details	19
3.5 Testing	20
3.5.1 Parameter Precision Results	21
3.5.2 Quantitative Transcoding Results	22
3.5.3 Visual Performance	26
3.6 JPEG 2000 to WSQ	27
3.7 Conclusions	28
4. Summary	29
Bibliography	31
Appendix A. Test Fingerprint Imagery	33
Appendix B. Image Quality Metric Data	35

Appendix C. NIST Code Modification	37
Glossary	39

List of Figures

Figure	Page
1. Mallat and WSQ Wavelet Decompositions	3
2. Processing Chain for Imagery	9
3. Average PSNR and IQM for Different Processing Chains	10
4. IQM on Individual Fingerprint Images	11
5. PSNR on Individual Fingerprint Images	11
6. Processing Chain for Tests	21
7. IQM Comparison of Transcoding Options	24

List of Tables

Table	Page
1. WSQ vs. JPEG 2000	6
2. Test Images	8
3. Actual Compression Ratios Achieved	10
4. Summary of JPEG 2000 Transcoding Settings	17
5. Subband Ordering	18
6. Number of Pixels Differing from Ground Truth	22
7. JPEG 2000 Standard Transcoding Decoded with $r=0.56$	23
8. JPEG 2000 Standard Transcoding Decoded with $r=0.5$	23
9. JPEG 2000 Alternate Transcoding Decoded with $r=0.5$	25
10. Comparison of WSQ and JPEG 2000 Compressed File Size	26

Section 1

Introduction

This report explores the degree of compatibility between the wavelet-based WSQ fingerprint compression standard and JPEG 2000 compression. The study was funded by the FBI with a view to anticipating the impact of future interactions between JPEG 2000 and WSQ. MITRE was uniquely placed to undertake this study due to our currently active support to the FBI and the JPEG 2000 development community. As an introduction, this section provides some background information about the two algorithms and MITRE's connections in these communities. To aid the reader in understanding the context of the remainder of the report, a brief summary of the project schedule closes this section.

1.1 Background on JPEG 2000

JPEG 2000 is a new ISO international standard for general wavelet-based image compression that is expected to be widely available in the future. The JPEG 2000 standard is being developed and written at the international level by ISO/IEC JTC1/SC29 WG1¹, informally known as the JPEG committee. This is the same committee that generated the current JPEG standard. Within each country, there are groups that participate in this effort. NCITS/L3.2² is the U.S. version of the JPEG committee.

Participants on the JPEG committee include commercial companies, universities, and other organizations from around the world. Many participants are active researchers in the field of image compression. The range of participants is broad: covering imagery production (camera, film, scanner, copier, printers, satellites), imagery software (database, browser, publication), government agencies, and contractors. Areas of interest include Internet, digital cameras, medical imaging, remote sensing, mobile applications, and motion video to name a few.

MITRE, represented by the author, has been a member of the JPEG committee for over four years, participating both nationally and internationally. By actively participating in meetings, MITRE has been able to provide a voice for our sponsors' concerns and stay abreast of the progress of JPEG 2000. MITRE submitted its own wavelet algorithm [1] during the JPEG 2000 call for contributions and since then has run numerous experiments jointly with other member companies to advance the development and understanding of JPEG 2000. Most recently, MITRE has actively participated in checking, clarifying, and correcting the technical content of the standards document.

The JPEG 2000 format is designed to allow very good compression of a wide variety of image types and has not been specifically tuned for fingerprint imagery. It is a decoder-only standard, in that it specifies how a file is decoded, but does not place restrictions on compression ratio, in order to serve a wide variety of applications with varying needs in the file size and image quality trade-off.

¹ International Organization for Standardization/International Electrotechnical Commission, Joint Technical Committee 1, Subcommittee 29, Working Group 1

² ANSI Accredited Standards Committee, National Committee for Information Technology Standards. L3.2 is the Still Image Coding working group

1.2 Background on WSQ and FBI Involvement

WSQ is the FBI specified fingerprint compression standard that is used throughout the United States criminal justice system and internationally.

MITRE's experience working with WSQ dates back to the early 1990s. This earlier work involved analysis of the WSQ compression technique and its incorporation into a simulation model of a large FBI system (IAFIS³). Also, experiments were run to investigate the effects of very high quality fingerprint scans on WSQ. The current investigation's base of WSQ knowledge relies on the WSQ Standards document [2], and utilizes the available WSQ source code for verification experiments within the context of our current image standards development support to the FBI.

Although the underlying WSQ format can be used to encode many image types, certain parameter settings have been specifically tuned to fingerprints. The FBI WSQ standard is both an encoder and decoder standard, and it places stringent restrictions on encoders, as well as decoders in order to maintain strict quality control.

1.3 Technical Background

This report assumes a certain amount of knowledge of the terminology used in the wavelet image compression field, and for in-depth reading an intimate knowledge of certain aspects of both WSQ and JPEG 2000. Due to limitations of project scope, time and space, definitions of compression terminology and an exact description of the two standards are not provided in this report. However, a glossary of summarized definitions is included and the bibliography mentions books, reports, and web-sites that provide further information.

1.4 Project Summary

This study consisted of two main segments. The first segment was a theoretical comparison between the two compression techniques, identifying where they are the same and where different, with a view towards identifying potential migration paths and coexistence paths. This study segment, reported in Section 2, was completed during June 2000, prior to the finalization of some aspects of JPEG 2000. In the second study segment, a testbed for converting WSQ files into JPEG 2000 files was created, tested, and analyzed. That study, reported in Section 3, was begun in October 2000 and completed in December 2000. During the interval between the two project segments, MITRE was at the forefront of actively promoting inclusion of WSQ-compatible elements into JPEG 2000. Our participation in the JPEG committee meetings required to forward these modifications was funded via a separate MITRE overhead project.

The final standardization of certain aspects of JPEG 2000 mentioned in this report will not occur until after this report is published, so the final JPEG 2000 standard documents should be inspected for ultimate verification of the concepts presented here.

³ Integrated Automated Fingerprint Identification System

Section 2

Interoperability Analysis

This section first summarizes some of the salient technical aspects of both standards [2,3], then focuses on the differences, and finally discusses consequences and possible avenues of future investigation. Only those technical aspects of the algorithms that impact interoperability and possible migration are discussed. Additional features of JPEG 2000 [3,4] are not addressed in this report.

2.1 WSQ Gray-Scale Fingerprint Image Compression Standard

Although WSQ is a format with a large degree of flexibility, many of the degrees of freedom have been eliminated by a fixed specification for the purposes of the Fingerprint Image Compression Standard. Therefore, for the remainder of this document, references to WSQ will refer to the fingerprint standard version of this algorithm only, unless otherwise specified.

The first stage of WSQ is to shift and scale the image data so that it is somewhat balanced around zero instead of being entirely positive. The shift and scale values are fixed-point floating values that are specified in the compressed bitstream.

WSQ currently has one approved wavelet filter, the Daubechies (9,7), applied with the one approved subband decomposition. This subband decomposition has a structure that is more complex than a simple Mallat or packet decomposition. In particular, at one of the resolution levels, the HH subband is decomposed differently than the HL and LH subbands, as shown in Figure 1.




Figure 1. Mallat and WSQ Wavelet Decompositions

After the wavelet transform, WSQ applies a scalar quantizer that uses a zero bin that is 1.2 times the step size of the non-zero bins. The relative step size for each subband is typically computed using energy level calculations that are performed on the transformed image data, but for a few subbands the data is quantized entirely to 0, and in a few others, there is no energy level scaling. The absolute step sizes, which also incorporate a scale factor related to desired compression rate, are encoded into the bitstream. The inverse quantizer in the decoder uses not only the transmitted step sizes, but also a global quantization bin center value that is also transmitted in the bitstream. For fingerprint WSQ the quantization bin center is always set to 0.44.

After quantization, transform coefficients are traversed in a specific order and entropy coded using a combination of Huffman and runlength coding. Only two Huffman tables are allowed in each fingerprint WSQ. Huffman tables are encoded into the bitstream.

Finally, all the encoded image data and associated parameters are combined in a syntax that uses markers and marker segments, similar in form to the current JPEG. The marker syntax is very explicit; data and parameters must appear in exactly this format in order for a bitstream to be decoded.

2.2 JPEG 2000

JPEG 2000 will be a standard that is issued in several parts. Part 1 will contain elements of the standard that any JPEG 2000 compliant decoder must understand. In order to facilitate implementation in many different environments, Part 1 encompasses a very restricted set of the potential capabilities of JPEG 2000. Extended optional functionality for still imagery will be contained in Part 2. Implementers are free to include any subset of Part 2 functionality in addition to the required Part 1 decoder functionality. Because there is a large difference in the capabilities of Part 1 and Part 2, this section will discuss them independently.

The Final Draft International Standard (FDIS) of JPEG 2000 Part 1 was approved in January 2001 and the International Standard of JPEG 2000 Part 1 (IS 15444-1) should follow shortly. The contents of Part 2 were still under review during this study and able to change until the committee draft release (August 2000). The final approval of Part 2 as an FDIS will not occur before July 2001.

2.2.1 Part 1

Image data is initially shifted so that it is somewhat balanced around zero instead of being entirely positive. The shift value is always a power of two, predefined based upon the dynamic range of the data, and is not signaled in the bitstream. No scaling is specified, though effectively power of two scaling can be contained in the implementation. Also, any scaling that applies to the full image can be applied across all image wavelet subbands by scaling the quantizer step sizes.

There are two wavelet filters allowed in Part 1, one of which is Daubechies (9,7). Only the Mallat subband decomposition is allowed in Part 1.

All data is quantized using a scalar quantizer with a zero-bin twice the size of the other bins. When using the (9,7) filter, the fixed-point floating step size values may vary from subband to subband. Choice of the step sizes is application-dependent and they are encoded in the bitstream. The inverse quantizer in the decoder uses not only the transmitted step sizes, but also a variable quantization reconstruction factor that the decoder is free to choose within the range [0,1). This factor is quite similar in intent to WSQ's quantization bin center, with the difference being that the reconstruction factor may vary during the decompression, and a specified value is not transmitted in the bitstream.

After quantization, transform coefficients are traversed by bitplanes in a specific order and entropy coded using arithmetic coding. Since the arithmetic coding naturally adapts to the data statistics, no tables need be transmitted for the entropy encoding. Once the bitplanes for individual subbands are encoded, they are placed into one of several progression orders. The progression order used is specified in the bitstream.

Finally, all the encoded image data and associated parameters are combined in a syntax that uses markers and marker segments, similar in form to the current JPEG but different enough from the WSQ syntax that they are not cross-interpretable.

2.2.2 Part 2

At the time of this study, JPEG 2000 Part 2 had not yet been completely specified, so the statements in this section are based upon an understanding of Part 2 as of June 2000. Since the specification was not yet complete, there was a possibility to have the committee adopt extra capabilities that would increase the compatibility between JPEG 2000 Part 2 and WSQ. Introducing such proposals would require code implementation and experimental support in a very short time frame.

As of June 2000, no changes in the original image shift structure had been included for Part 2. However, the MITRE representative to the JPEG committee initiated preliminary discussion of a more flexible shifting structure and received a generally positive response.

More general transforms will be allowed in Part 2, though there may be limitations based on filter length and other factors. More generalized subband decompositions will be allowed in Part 2. However, as of June 2000, the flexibility of these decompositions was still to be determined. In particular, the decomposition used by WSQ was not allowed by Part 2, but several companies wanted to remove this restriction.

Other quantizers are envisioned for Part 2. In particular, the Trellis Coded Quantizer (TCQ) will be included as a Part 2 option. The quantizer specified in WSQ had not yet been integrated as a Part 2 option in June 2000. However, when MITRE initiated preliminary discussion of a more flexible scalar quantizer for Part 2 at JPEG meetings, the feedback was generally positive.

No alternate transform scanning orders are envisioned for Part 2. In spite of initial discussion of alternate entropy encoders, there is very little support for this idea.

MITRE asked about alternate syntaxes at a JPEG meeting and received a very negative response from the attendees queried.

2.3 Evaluation

Table 1 summarizes the information in the previous sections. Areas where JPEG 2000 can mimic WSQ are unshaded, while conflicting segments are shaded in dark gray. Lightly shaded areas are proposals defined in June 2000, which might allow JPEG 2000 Part 2 to mimic WSQ.

As can be seen in Table 1, there are quite a few places where the two algorithms differ. Although many of these differences are minor, in that they perform quite similar operations, they do cause considerable differences between a WSQ-compressed bitstream and a JPEG 2000-compressed bitstream.

Depending upon where differences occur between compression algorithms, there are a variety of methods that can be used to change between two formats. We describe three possible techniques.

- 1) At one extreme, when the algorithms are entirely different (particularly at the first encoder stages), a compressed file must be totally decompressed to an image with the appropriate decoder, and then

Table 1. WSQ vs. JPEG 2000

Algorithm Function	WSQ	JPEG 2000 – Part 1	JPEG 2000 – Part 2	
			June 2000	Proposals
Image Offset	Contained in bitstream. Float value allowed.	Cannot be specified. Predefined values used. (shift of 128 for 8-bit images)		Add general shift.
Image Scale	Contained in bitstream. Float value allowed.	Can be specified by adjusting quantizer step size.		
Wavelet Filter	Daubechies (9,7)	Daubechies (9,7) and an Integer (5,3)	Daubechies (9,7) one of many	
Wavelet Decomposition	Special customized tree for fingerprints	Mallat tree only	More general trees, but not WSQ tree.	Include WSQ tree.
Frequency Weighting	Scaling: Subband data dependent formulas.	Scaling: May be set by subband. Encoder only detail. Could mimic WSQ.		
Quantizer	Scalar: 1.2Q zero bin	Scalar: 2Q zero bin	Add TCQ	Add generalized scalar quantizer.
Inverse Quantizer Reconstruction	Scalar bin center (0.44) sent in bitstream	Scalar bin center: may be chosen at will by decoder.		
Scan Order	Raster within subbands	Vertical stripe scan		
Entropy Coding	Huffman + runlength	Arithmetic with runlength		
Bitstream Ordering	Fixed order: progressive by resolution	User selectable: Progressive by resolution is one option.		
Syntax	Modified JPEG	Different modification of JPEG	Extended version of JPEG 2000 Part 1	

the new image is recompressed with the other compression algorithm. Although this is always doable, there is a downside in that the image quality may suffer in the process.

2) The opposite extreme occurs when the two algorithms, including syntaxes, are so similar that one is a subset of the other. In that case, it may be possible to have the decoder from one algorithm directly read a compressed file from the other algorithm. In practice, this is extremely unlikely to happen unless the new algorithm is created with this property as a specific design goal. WSQ and JPEG 2000 do not have this property.

3) Between the two extremes is a range of possibilities that all include some amount of partial decoding with one algorithm, followed by partial encoding with the other. For instance, if two algorithms differ only in the entropy coder and syntax, then one can decode one syntax and entropy coder to the level of the quantized coefficients. The data can then be recoded using the other entropy

coder and syntax. This process is referred to as transcoding. If the transcoding does not affect the quantization or other lower level functionality, then there is no impact on image quality as a result of the transcoding. If, however, the quantization or lower level processes are affected, then there is likely to be some impact on image quality.

From Table 1, we see that WSQ cannot be transcoded to JPEG 2000 Part 1 when the image offset is different from 128. Since there is no guarantee that the image offset in WSQ will be 128, it will be necessary to entirely decode WSQ data before recompressing it with JPEG 2000 Part 1. Likewise, since JPEG 2000 Part 1 uses a different wavelet decomposition tree than WSQ, any transcoding from JPEG 2000 Part 1 to WSQ will require changes in the quantization and additional sections of wavelet decomposition. This can have a slight negative impact on image quality.

Since JPEG 2000 Part 2 was not completely specified during this study, some of the segments needed to facilitate transcoding might be added as options. Useful options would be control of image offset, WSQ wavelet decomposition, and scalar quantizer with 1.2Q zero bin. MITRE's initial discussions with a few members of the JPEG committee showed some support for these three additions. Although perfect compatibility between the algorithms would be useful to the forensics and law enforcement community, it is unlikely that the entire WSQ process will be adopted into JPEG 2000 Part 2. In particular, elements such as scan order and syntax are very unlikely to come into agreement.

2.4 Compression Performance

To get an idea of the differences in compression performance between WSQ and JPEG 2000 and to see what loss in quality might occur due to image recompression, an initial test was run on a small set of fingerprints. As shown in Table 2, this set consisted of 12 rollprint, right index finger images, selected from a cross-section of FBI-certified card scanners and latent scanners, with light, medium, and dark inked fingerprints, and several live scan images. Since the available Verification Model (VM7.0) for JPEG 2000 did not contain the three proposals that would enable JPEG 2000 to be equivalent to WSQ through the quantizer stage (allowing no-loss transcoding), this comparison only includes WSQ and JPEG 2000 Part 1.

Test procedure:

- 1) Compress origimage with "wsq_demo", an FBI-certified version of WSQ (WSQ by Aware©, Inc., Version 1.73, dated 6 November 1995) using '-ratio 12'. We found that this input parameter gave an average effective compression rate near 15:1.
- 2) Compress origimage with VM7.0 to match as closely as possible the effective compression rate achieved by wsq_demo and such that the VM7.0 file is never larger than the wsq data file. (See Table 3 for actual compression ratios.) Default settings were used by VM7.0, except for the following flags: -step 0.003956 -rate <rate>
- 3) Decompress the two files to produce a wsqimage and a j2kimage.
- 4) Using the same settings as in the first 2 steps, recompress wsqimage and j2kimage with both wsq_demo and VM7.0. Decompress these files as well.

Output images: For each original fingerprint image, six reconstructed images were generated corresponding to differing processing paths as shown in Figure 2.

Table 2. Test Images

Image	Source	PPI	Size	Comment
D1_125	DBA <i>Umax PowerLook III</i>	1000	1102 x 1426	Light-inked, card #29
M1_15	Mentalix <i>Umax PowerLook III</i>	1000	1073 x 1275	Light-inked, card #15
D1_377	DBA <i>Umax PowerLook III</i>	1000	1360 x 1348	Medium-inked, card #57
M1_92	Mentalix <i>Umax PowerLook III</i>	1000	1140 x 1480	Dark-inked, card #92
D5_125	DBA <i>Umax PowerLook III</i>	500	583 x 715	Light-inked, card #29
D5_235	DBA <i>Umax PowerLook III</i>	500	613 x 533	Light-inked, card #27
D5_377	DBA <i>Umax PowerLook III</i>	500	681 x 695	Medium-inked, card #57
D5_582	DBA <i>Umax PowerLook III</i>	500	512 x 704	Dark-inked, card #92
H_WJ	HBS <i>LSI/T+</i>	500	533 x 719	livescan, light impression
H_KO	HBS <i>LSI/T+</i>	500	512 x 735	livescan, normal impression
X_H2	CrossMatch <i>ID1000</i>	500	595 x 652	livescan, normal impression
X_A2	CrossMatch <i>ID1000</i>	500	579 x 681	livescan, normal impression

Notes: All images are right index finger rolls, acquired on FBI-certified scanners.
Ten-print cards are from the FBI's Fingerprint Card Master File test set.
Life size views of all these images are shown in Appendix A.

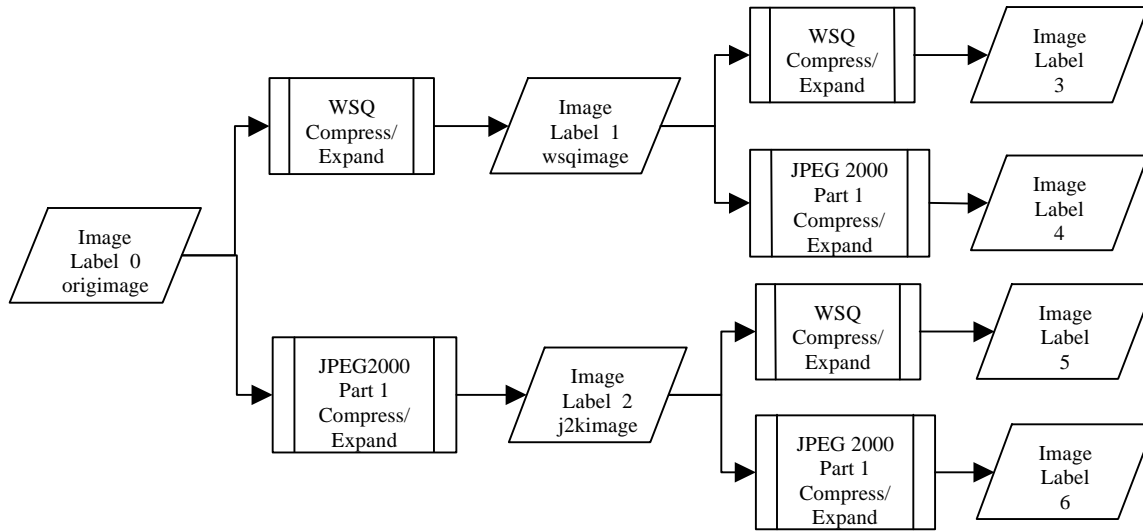


Figure 2. Processing Chain for Imagery

2.4.1 Image Metrics

Two image quality metrics were applied to the data: PSNR and IQM. PSNR (peak signal-to-noise ratio) [5] is a commonly used metric for indicating how faithfully one image matches another (original) image. It is well known that when there are small changes in image quality, PSNR will often conflict with visual testing results, so care must be used when interpreting PSNR results. IQM (Image Quality Metric) is a metric developed by MITRE [6] to be used as an absolute quality indicator with no need to compare against an original image, designed to more closely approximate visual test results.

PSNR was computed for each reconstructed image relative to the original over the entire image area, while IQM⁴ was computed separately for all the images (original and reconstructed) in a restricted image area. The two metrics are self-consistent but in direct conflict with each other, so visual tests are a necessity. Raw PSNR and IQM values are tabulated in Appendix B.

Figure 3 shows the average PSNR and IQM results from all the test images for each of the reconstructed image types. This shows some general data trends between the processing paths.

Figures 4 and 5 compare the results from just a single compression pass with the two algorithms separately on an individual image basis. The images are grouped so that the first four are 1000 ppi inked card scans, the next four are 500 ppi inked card scans (somewhat lower absolute IQM), and the last four are 500 ppi livescans (somewhat higher IQM than the 500 ppi inked cards). Within each group of four, the ordering is from lightest scan to darkest. The IQM tends to increase from light to normal inking but eventually decreases if the inking becomes too dark. Since PSNR is a relative image metric, it does not show these differences.

⁴ IQM version 5.5 was run with sensor 5 option and with the suboption to take into account resolution level (1000 ppi versus 500 ppi); 512x512 window for 500 ppi images, 1024x1024 window for 1000 ppi images.

Table 3. Actual Compression Ratios Achieved

Image	WSQ	J2	WSQ ↓ WSQ	WSQ ↓ J2	J2 ↓ WSQ	J2 ↓ J2
D1_125	11.85	11.87	11.94	11.87	10.99	11.85
M1_15	15.13	15.14	15.16	15.16	15.18	15.13
D1_377	13.50	13.52	13.51	13.52	13.19	13.50
M1_92	13.34	13.34	13.40	13.35	12.44	13.34
D5_125	12.00	12.02	12.05	12.00	11.54	12.00
D5_235	16.19	16.33	16.32	16.35	14.76	16.19
D5_377	13.95	14.00	13.96	13.95	12.86	13.95
D5_582	14.53	14.70	14.55	14.54	13.94	14.55
H_WJ	19.02	19.13	18.98	19.15	17.88	19.03
H_KO	16.17	16.22	16.16	16.21	14.35	16.18
X_H2	15.71	15.73	15.70	15.85	14.53	15.71
X_A2	17.17	17.28	17.19	17.20	15.63	17.19




Figure 3. Average PSNR and IQM for Different Processing Chains




Figure 4. IQM on Individual Fingerprint Images

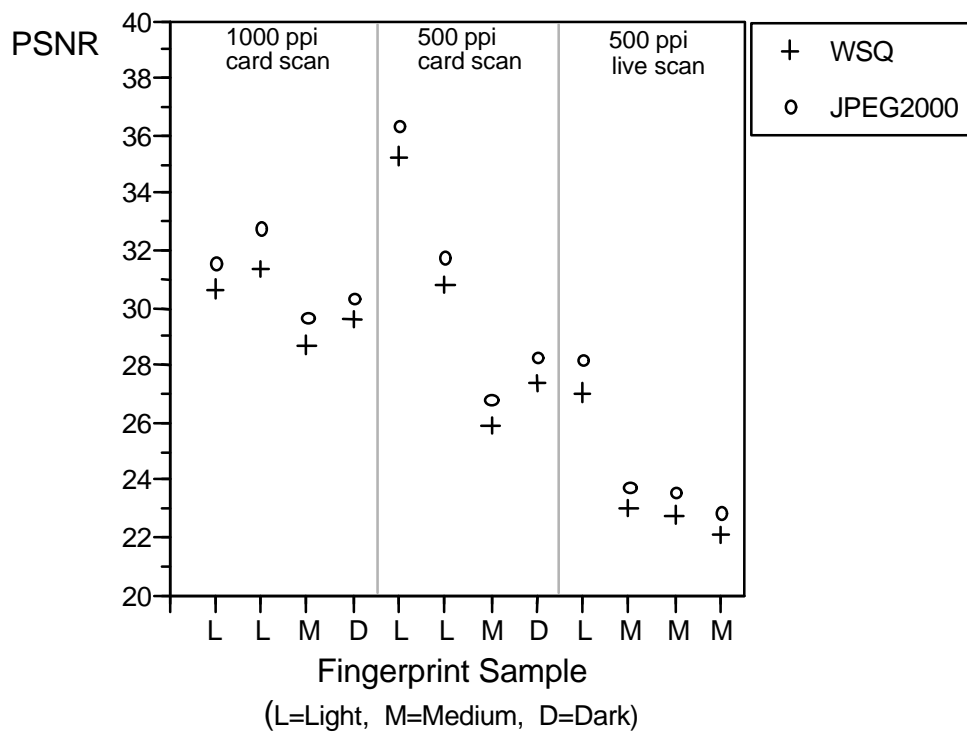


Figure 5. PSNR on Individual Fingerprint Images

2.4.2 Visual Performance

Two observers performed the visual inspection. Both have image processing/evaluation experience, including previous work with fingerprints, although they are not formally trained in fingerprint identification or matching. When two images were very similar, a flicker was used to identify differences. (Flickering rapidly between two images created movement in areas where the images differ, and is much more sensitive to differences than side-by-side comparisons.) The comments below summarize the findings of this comparison.

The reconstructed images from WSQ and JPEG 2000 Part 1 are very similar at first glance, but with close inspection a few differences are noted. JPEG 2000 Part 1 has a slightly softer appearance, while WSQ looks slightly crisper and seems to approximate very fine texture better.

However, the retention of ridge, bifurcation, and sweat pore information seems very similar in both. How these slight differences might impact human and machine fingerprint matching results is unknown.

For the 1000 ppi scans, the changes caused by the different processing paths were smaller than the size of any fingerprint features. Differences could only be perceived when two images were flickered rapidly, and the changes appeared to be at the noise level.

For the 500 ppi scans, the image X_A2 showed the differences between the algorithms most prominently, but even in this image in some areas WSQ preserved features better, and in others JPEG 2000 Part 1 had the edge. On the other images, it was much more difficult to identify differences between the algorithms. This may be in part due to the fact that the compression ratio used for X_A2 was somewhat higher than normal. H_WJ also has a very high compression ratio, but since the print is so light it is much harder to see differences between the algorithms on it.

The processing paths that repeated the same compression twice in a row were visually identical to processing with that algorithm only a single time (flicker nearly imperceptible). The paths that used the two different compressions in sequence had slight changes from either WSQ or JPEG 2000 Part 1 alone, and seemed to combine features of both.

2.5 Comments on VM7 vs. WSQ

VM7 is one particular encoder that is being used to test JPEG 2000 concepts. It is not currently set up to do some of the extra processing used in WSQ, but there is nothing in the JPEG 2000 standard that would prevent some customization. The following are some places where VM7: JPEG 2000 Part 1 could be further customized to be more similar to WSQ. Any of these modifications might slightly change the JPEG 2000 Part 1 results shown in this section, but none of them will substantially change the image quality.

- WSQ specifies that the quantization step sizes should be computed based upon statistics from a restricted area of each subband. Since JPEG 2000 is allowed to specify any step sizes, it would be possible to do this on the subbands in Part 1. However, since the VM7 code does not presently do this, we have not been able to see what improvements this might generate in the results.

- WSQ also uses weighting of the subbands. Since JPEG 2000 Part 1 cannot compute the WSQ subband decomposition it is difficult to duplicate this. However, the very highest frequency diagonal bands are zeroed out in WSQ, and it would be possible to force VM7 to do this.
- Also the VM7 code is set up to currently use a bin center reconstruction factor of 0.5, and the encoder makes some decisions based upon the assumption that 0.5 is being used. Using 0.44 as in WSQ would require some alterations in the code and would alter the results slightly.

2.6 Conclusions

JPEG 2000 Part 1 general COTS compression products when used on fingerprints will create a file that is slightly lower quality than WSQ at the same file size. Tuning of the JPEG 2000 Part 1 parameters to fingerprints will improve performance somewhat, but we suspect there will always be a slight gap due to the difference in decomposition structure.

The algorithmic comparison in this section shows that it would be much easier to move between WSQ and JPEG 2000 formats if a few WSQ elements were incorporated into JPEG 2000-Part 2. Proposing additions within committee requires implementing the concepts in the current Verification Model code, generating text for the standards document, and having some experimental evidence that these elements improve or at least do not degrade transcoding performance.

Given initial positive response to the idea, there was a reasonable chance that extensions to the image shift, wavelet decomposition, and scalar quantizer would be adopted, if code and text were already available by 3 July 2000. MITRE generated VM7 code, text, and test results for the level shift and generalized scalar quantizer during June 2000. Concurrently, SAIC implemented more generalized wavelet decompositions. After MITRE's presentation of the test results and discussion within the JPEG committee, all three proposed additions were formally accepted into Part 2.

Section 3

Transcoding

The previous section examined the theoretical differences between WSQ and JPEG 2000 and determined that several new features were required in JPEG 2000 Part 2 to enable transcoding. After that initial study the three elements needed to allow increased compatibility with WSQ were formally incorporated as options for the JPEG 2000 Part 2 standard. The next study incorporated these new elements into a testbed for converting data from WSQ format to JPEG 2000 Part 2 format. This section details how compressed data is converted from one format to the other and documents and evaluates differences that appear during the conversion process.

The section begins with technical aspects of the format conversion and continues with a discussion of elements that may cause loss of precision in the conversion. Then testbed implementation details are summarized, followed by testing details and actual results on fingerprint images. The section finishes with some general conclusions about WSQ/JPEG 2000 transcoding.

3.1 Notation

A WSQ file is comprised of a number of control parameters and Huffman encoded quantized transform coefficients. The variable parameters that impact the conversion are:

R = image data scaling
M_w = image data shift
Q_w = quantization bin size for each subband
Z = zero bin size for each subband

In addition, the (9,7) wavelet transform filter coefficients (with $\sqrt{2}$ normalization) and a fixed reconstruction bin center $C=0.44$ are transmitted. A fixed wavelet decomposition is used with the subband ordering specified in [2]. There are also other parameters passed concerning the Huffman encoding, but they have no affect on the conversion process.

A JPEG 2000 file also has a number of control parameters. These include:

M_J = image data shift (DC offset)
Q_J = quantization bin size for each subband
NZ = shrinkage in size of the zero bin for each subband
GenDecomp = information on the wavelet decomposition tree

One of the default filters allowed is the (9,7) filter with (1,2) normalization, in which case no filter coefficients are transmitted. The reconstruction factor ' r ' has a function very much like the reconstruction bin center C in WSQ, but is not specified by the standard and may be adjusted as desired by the decoder. A different subband ordering is used in JPEG 2000, as specified in [7].

3.2 Conversion

To perform the conversion, the WSQ notation must be transformed into a form that matches the JPEG 2000 notation. This involves addressing wavelet filter normalization, wavelet decomposition

specification, image shifting/scaling, setting quantization bin sizes, and choosing the reconstruction factor.

Tree Decomposition: The particular decomposition used in WSQ must be specified within the JPEG 2000 Part 2 framework [7]. This is done by specifying:

Number of decomposition levels:	$N_L = 5$
Resolution level structure:	$I_R = 0$ (i.e. default joint split at each level)
Decomposition depth at each level:	$I_\theta = 4$ $d_\theta = 2321$
Sublevel splitting structure:	$I_S = 17$ $d_S = 0110111111111111$

This decomposition structure gives the same split as WSQ *except* that the last 4 subbands (60-63) are joined into one band. Since these bands are always quantized to zero in the FBI's WSQ standard and never transmitted, this combination of the four subbands into one, which is also quantized to zero, makes no change in the reconstruction. (If the user desires the exact same split as WSQ, then the first 0 in the sublevel splitting structure shown above can be changed to a 1.)

At a purely implementation level, the difference in subband ordering must also be handled. Since it is somewhat complicated the exact details of this relationship are provided separately in section 3.2.1.

Wavelet Filter: The irreversible (9,7) wavelet filter specified in JPEG 2000 Part 1 is a scaled version of the wavelet specified in the FBI's WSQ standard. If this filter is used, there is no need to transmit the wavelet filter coefficients. The (9,7)-wavelet filtering operations used by WSQ and JPEG 2000 differ primarily in the normalization that is used. The WSQ filter has magnitude gains of $(\sqrt{2}, \sqrt{2})$ for low-pass and high-pass, while the JPEG 2000 filter typically has gains of (1, 2) or (1,1) depending upon the implementation and interpretation of bin sizes. (This choice is an implementation issue only. Step sizes are reported in the compressed bit-stream relative to the implementation gain, and can be correctly interpreted by implementations using the other filter normalization.) For the sake of simplicity this analysis assumes the (1,1) normalization used in the VM8.5⁵ JPEG 2000 implementation.

This difference in normalization means that in order to generate JPEG 2000 wavelet coefficients the WSQ wavelet coefficients must be divided by a gain of 2 for every level of two-dimensional (2-d) transform applied.

$$T'(x) = T(x) / 2^n \quad \begin{array}{l} T \text{ and } T' \text{ are the WSQ and JPEG 2000 transforms respectively.} \\ n = \text{number of 2-d decompositions required to obtain a subband.} \end{array}$$

Image Scaling: The image scaling applied in WSQ has no exact correlate within JPEG 2000. However, since scaling image data prior to the wavelet transform is equivalent to scaling wavelet coefficients after the transform, i.e.

$$T'\left(\frac{I-M}{R}\right) = \frac{T'(I-M)}{R},$$

it is possible to incorporate the scaling factor into the quantization operation.

⁵ VM8.5 was the JPEG 2000 test code used within the JPEG committee during October-November 2000.

If $Q' = RQ$ and $Z' = RZ = 1.2 Q'$, we see

$$\frac{T'(x/R) - Z/2}{Q} = \frac{T'(x)/R - Z/2}{Q} = \frac{T'(x) - RZ/2}{RQ} = \frac{T'(x) - Z'/2}{Q'}$$

Quantization Step Sizes: By incorporating the effects of gain and image scaling into the quantization step size within JPEG 2000, it is possible to obtain results that are theoretically equivalent to WSQ. The JPEG 2000 step sizes assuming an implementation that uses the (1,1) normalized (9,7) filter are:

$$Q_J = Q_W R / 2^n \quad \text{where } n = \text{number of 2-d decompositions performed to obtain the subband}$$

In addition, it is necessary within VM8.5 to specifically mark any bands that will be quantized to zero, since the step sizes within JPEG 2000 are not allowed to be either 0 or infinity. This allows the entire subband to be encoded as zero, and assigns a legal but arbitrary step size.

Zero Bin Size: The zero bin in WSQ is specified as 1.2 times the size of the regular bin size, but this value is also written into the WSQ file, and due to the inability of a binary system to perfectly represent this value, the actual value transmitted fluctuates slightly. In JPEG 2000 the default zero bin width is two times the regular step size. However it is possible, via a Part 2 option to specify shrinkage in the default zero bin size. The best calculation for shrinkage is $NZ = 1 - \frac{1}{2} Z/Q_W$. This produces $NZ \approx 0.4$ when converting WSQ to JPEG 2000.

Image Shift: The image shift 'M' is identical in both standards and is just translated from one file format to the other. $M_J = M_W$. When this value is anything other than 128, the JPEG 2000 Part 2 DC offset capability must be enabled.

Reconstruction Factor: The choice of reconstruction factor (bin center) within JPEG 2000 cannot be indicated in the compressed data. However, if the decoder knows it is decoding transcoded fingerprint data, then it has the option to set the reconstruction factor to mimic $C=0.44$. Although there are differences in terminology between WSQ and JPEG 2000, it can be shown that $r = 1-C$. So a value of $r=0.56$ will mimic a WSQ decoder. (How a decoder would know it is decoding a WSQ transcoded file would be up to the user community and/or implementation. User defined tags can be generated for use within JPEG 2000 and could be used for this purpose.)

Table 4 summarizes transcoding settings recommended in this section.

Table 4. Summary of JPEG 2000 Transcoding Settings

Transform	Filter: default (9,7) irreversible
	$N_L = 5 \quad I_R = 0$
	$I_\theta = 4 \quad d_\theta = 2321$
	$I_S = 17 \quad d_S = 0110111111111111$
	Subband Ordering: see Table 5
DC Offset	$M_J = M_W$
Quantization	Assuming (1,1) filter normalization, $Q_J = Q_W R / 2^n$ where $n = \#$ 2-d decomp
	$NZ = 1 - \frac{1}{2} Z/Q_W$
Decoder Option	$r = 0.56$

3.2.1 Subband Ordering

JPEG 2000 uses an ordering of subbands, $O(\cdot)$, that is different from the frequency weighted ordering used in WSQ. Table 5 relates the WSQ order with the JPEG 2000 order. This subband reordering must be taken into account when setting Q_j and NZ and when transferring the Huffman decoded data from WSQ to JPEG 2000.

Table 5. Subband Ordering

WSQ	JPEG 2000			Sequence of 2-d transforms (0=LL,1=HL,2=LH,3=HH) First to last from left to right. n =number of entries		JPEG 2000			WSQ
Cnt	$O(\cdot)$	Band Index ⁶	Lev			Lev	Band Index	$O(\cdot)$	Cnt
0	0	0	0	0,0,0,0,0		4	27	30	31
1	1	1	1	0,0,0,0,1		4	26	29	32
2	2	2	1	0,0,0,0,2		4	25	28	33
3	3	3	1	0,0,0,0,3		4	24	27	34
4	4	1	2	0,0,0,1		4	40	43	35
5	5	2	2	0,0,0,2		4	41	44	36
6	6	3	2	0,0,0,3		4	42	45	37
7	8	5	3	0,0,1,1		4	43	46	38
8	7	4	3	0,0,1,0		4	45	48	39
9	10	7	3	0,0,1,3		4	44	47	40
10	9	6	3	0,0,1,2		4	47	50	41
11	13	10	3	0,0,2,2		4	46	49	42
12	14	11	3	0,0,2,3		4	34	37	43
13	11	8	3	0,0,2,0		4	35	38	44
14	12	9	3	0,0,2,1		4	32	35	45
15	18	15	3	0,0,3,3		4	33	36	46
16	17	14	3	0,0,3,2		4	39	42	47
17	16	13	3	0,0,3,1		4	38	41	48
18	15	12	3	0,0,3,0		4	37	40	49
19	23	20	4	0,1,1,0		4	36	39	50
20	24	21	4	0,1,1,1		4	48	51	51
21	25	22	4	0,1,1,2		5	5	53	52
22	26	23	4	0,1,1,3		5	4	52	53
23	20	17	4	0,1,0,1		5	7	55	54
24	19	16	4	0,1,0,0		5	6	54	55
25	22	19	4	0,1,0,3		5	10	58	56
26	21	18	4	0,1,0,2		5	11	59	57
27	33	30	4	0,1,3,2		5	8	56	58
28	34	31	4	0,1,3,3		5	9	57	59
29	31	28	4	0,1,3,0		5	12	60	60-63
30	4	29	32	0,1,3,1					

⁶ Precise band index values are implementation dependent, but the ordering of these values will always be consistent with the values given here. These particular band index values are used in VM8.5.

3.3 Parameter Format/Precision

In addition to computing the parameter values for JPEG 2000, it is necessary to format them for transmission within the JPEG 2000 file format. Since different formats are used in WSQ and JPEG 2000 there can be a loss in precision at this stage. The object of this section is to examine these differences and examine the amount of error this may produce.

For the values, M_w , R , C , Q_w , and Z , WSQ uses the format $m10^{-e}$ where m is an integer of 16 bits and e is an unsigned integer. Therefore everything is in base 10.

JPEG 2000 uses different formats for different data types, but they are generally represented with an exponent base 2. In particular,

$$Q_j = (1 + \frac{m}{2^{11}})2^t \quad \text{where } \mu \text{ is an 11-bit unsigned integer and } \tau \text{ is a signed integer}$$

$$M_j = \beta/2^{16} \quad \text{where } \beta \text{ is a signed 32 bit integer}$$

$$NZ = v/2^{15} \quad \text{where } v \text{ is a signed 16 bit integer}$$

This means that values represented exactly in one format will not be able to translate exactly into the other. For example, the value 0.4 cannot be represented exactly in base 2, and since this is the value of NZ there will necessarily be some errors introduced at this point.

Also it is clear that since m has 16-bit precision and μ only 11, there is a loss of up to 4 bits in precision of Q_j in addition to any errors incurred representing the value base 2 rather than base 10. After the inverse transform, but prior to the final rounding to byte data, this loss in precision is smaller than one gray level. However, due to the non-linearity of the final rounding to unsigned byte image data, there are situations where a pixel will round to a value either above or below the value expected using the 16-bit precision. Examples of this rounding error can be seen in the test results.

The image shift/offset M is applied at the very last stage of reconstruction, just prior to rounding the floating point image data to the closest unsigned byte. Since M_j has 16 fractional bits of precision slight differences between M_j and M_w are unlikely to cause many differences between reconstructed images. (Note: Since the testing JPEG 2000 format for M_j has changed to 32-bit floating point.)

3.4 Implementation Details

The testbed was generated using NIST WSQ decompression source code⁷ combined with the VM8.5 JPEG 2000 compression source code⁸. A few bug fixes were needed in both algorithms. For the benefit of any readers who may be using the NIST code, the bugs found and fixed are discussed in detail in Appendix C. Bugs found in the VM will be fixed in future releases.

⁷ NIST wsq_v3_1 dated 2-13-95 available at <ftp://sequoyah.nist.gov/pub/src>. This code is not FBI-certified for WSQ. With the modifications applied, however, the test reconstructed images are within certification guidelines.

⁸ VM8.5 contains the extra elements required for WSQ transcoding. Since JPEG2000 Part 2 has not yet become an international standard, files generated by this code may differ in small ways from the final standard.

Marker segments prior to the actual Huffman encoded subband data were read using the NIST decompressor to extract image dimensions and the WSQ control parameters. This information was then used to initialize the JPEG 2000 compressor with conversions for subband ordering and the appropriate parameter calculations as summarized in Table 4. In addition, the JPEG 2000 wavelet filter and tree decomposition parameters were set as described in Table 4.

Since all data was to be transcoded with no attempts at embedding or rate control, the no_truncate flag was enabled within VM8.5.

Once the JPEG 2000 parameterization and initialization was complete, the NIST decoder began processing the Huffman encoded data. Each subband of quantized coefficients was Huffman decoded and then input to the VM8.5 fixed-point quantizer on a line-by-line basis. The VM8.5 fixed-point quantizer shifted the data into the position expected by the encoding process, but otherwise did not change the quantized transform coefficient value produced by the NIST decoder. All further processing within the VM8.5 compression then proceeded as normal.

This combined code was able to read a WSQ file and generate a decodable JPEG 2000 Part 2 file.

3.5 Testing

Since some differences in reconstruction are expected due to various changes in parameter precision, an initial test was performed to isolate effects of the changes in parameter formatting. This test used the NIST decompression source exclusively, enabling a flag that forced various decoded parameters to be changed into JPEG 2000 format. The number and type of differences in the reconstructed images was then recorded.

The end-to-end test then performed an actual transcoding from WSQ to JPEG 2000 Part 2 and made similar comparisons.

All tests were performed using compressed image files (name.wsq) and ground truth reconstructed image files (name.rec). A ground truth reconstruction is the most accurate reconstruction possible from the WSQ compressed file. For most of the tests, ground truth reconstruction was the reconstruction specified in the FBI's WSQ certification reference test set.

- 1) Decompress name.wsq with original NIST decompression to generate name.nist0. Also decompress name.wsq with NIST using flags to mimic JPEG 2000 parameter formatting. These reconstructions are called name.nist1-4.
- 2) Transcode name.wsq to name.j2k using the newly implemented testbed. Record compressed file sizes.
- 3) Decompress name.j2k using VM8.5 decompressor. Produces name.vm.
- 4) Compare reconstructed images (name.vm and name.nist<n>) with the ground truth reconstruction.

Except where indicated test data came from the WSQ certification reference test set. All of the *.wsq and *.rec files at ftp://sequoyah.nist.gov/pub/cmp_imgs/cmp_imgs/75 were used.

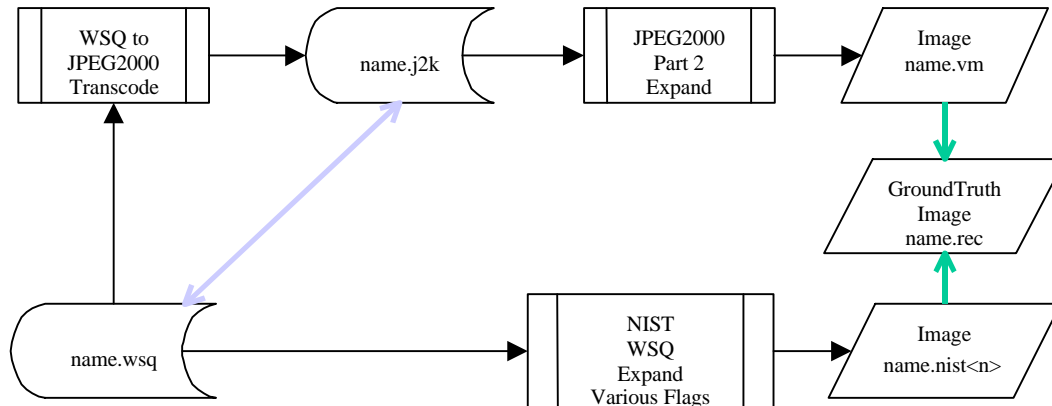


Figure 6. Processing Chain for Tests

3.5.1 Parameter Precision Results

Modifications of the NIST decoder were used to test the impact of the differing parameter formats, both in terms of restricted precision and changes from base 10 to base 2.

The NIST decoder was run in 5 different modes:

- 0) Use parameters as decoded using WSQ specifications.
- 1) (9,7) filter coefficients changed to more closely match JPEG 2000 implementation.
- 2) M_w was reformatted as M_j .
- 3) Z was reformatted as $2(1-NZ)Q_w$ where NZ was formatted as in JPEG 2000.
- 4) RQ_w was decoded as in JPEG 2000 format and replaced Q_w . R was set to 1.0 for later computations.
- 5) Combo of 1-4 above.

In all these cases, the error was never larger than one gray-level, when compared to the ground truth reconstruction `name.rec`. Moreover, the decoded image for mode 2 was always identical to the standard NIST mode 0; though both were slightly different than the ground truth. This indicates that, as anticipated, there is no loss in precision due to reformatting of M . However, there are small error contributions due to wavelet implementation and NZ formatting, and a much larger error contribution from the Q_j formatting.

Table 6 shows the number of pixels that differ from the ground truth reconstruction in each case. Obviously, the primary contributor to the error is the reduced precision for the quantizer step size in JPEG 2000. While the recorded error always affects less than one percent of the image area, it should be noted that this is beyond the 0.1 percent error tolerance specified for FBI-certified WSQ reconstructed values. It should also be noted that most of this error occurs in the fingerprint portion of the image rather than the background. A small check showed that these pixel differences affect less than 1.5 percent of the fingerprint area.

Table 6. Number of Pixels Differing from Ground Truth

Name	#Pix_Total	Mode0	Mode1	Mode2	Mode3	Mode4	Mode5 (Combo)	Combo %ImageArea
cmp00001	356345	4	6	4	20	1380	1458	0.4%
cmp00002	638976	7	26	7	23	1773	1865	0.3%
cmp00003	638976	0	8	0	6	1122	1143	0.2%
cmp00004	612880	5	20	5	16	2413	2417	0.4%
cmp00005	638976	2	23	2	11	4431	4420	0.7%
cmp00006	638976	3	26	3	21	5128	5176	0.8%
cmp00007	347710	5	10	5	17	1227	1294	0.4%
cmp00008	600000	8	24	8	32	3782	3866	0.6%
cmp00009	347136	5	11	5	20	945	1062	0.3%
cmp00010	197250	2	10	2	14	1031	1080	0.5%
cmp00011	440238	4	13	4	22	1145	1229	0.3%
cmp00012	369456	1	6	1	21	2971	3038	0.8%
cmp00013	350889	1	11	1	19	877	986	0.3%
cmp00014	269348	1	6	1	16	1496	1515	0.6%
cmp00015	292120	0	3	0	6	589	621	0.2%
cmp00016	504828	7	15	7	18	752	797	0.2%
cmp00017	346986	3	10	3	23	1199	1247	0.4%

3.5.2 Quantitative Transcoding Results

When the compressed WSQ files were actually transcoded, the reconstructed results were similar to those seen in the precision test. Since the JPEG 2000 decoder has a degree of flexibility in how the reconstruction value r may be set, a few different results are presented.

Standard Transcoding

In this test, files were transcoded using the formulas described in Table 4 and then decoded using two different decoder reconstructions.

The first test assumed that the decoder was customized for WSQ transcoded fingerprints, so the reconstruction factor $r=0.56$ was used. In this case, there was a maximum difference of one gray-level at any pixel. Table 7 shows the number and percentage of pixels where this error occurred. Although not identical to the numbers in Table 6, they are quite similar in magnitude.

The JPEG 2000 decoder, however, is not required to use $r=0.56$ and most JPEG 2000 decoders are likely to use $r=0.5$, or something smaller when given no other direction. Therefore, a second test decoded the transcoded files using a more generic JPEG 2000 decoder with $r=0.5$. In this case, the differences from ground truth became larger than one gray-level, and many more pixels had small differences from the ground truth reconstructed value. Table 8 shows the results of this experiment.

Since a large proportion of the pixels have some small variation from the ground truth, a further metric test was applied, namely IQM. To allow comparisons to the previous IQM results and include some livescan imagery in the mix, the imagery set described in Section 2 was used in the IQM test.

Table 7. JPEG 2000 Standard Transcoding Decoded with $r=0.56$.
Number of pixels off ground truth by one gray-level.

Name	#Pix_Total	JPEG 2000 $r=0.56$	%PixError
cmp00001	356345	1093	0.3%
cmp00002	638976	2613	0.4%
cmp00003	638976	2157	0.3%
cmp00004	612880	699	0.1%
cmp00005	638976	1696	0.3%
cmp00006	638976	4819	0.8%
cmp00007	347710	1654	0.5%
cmp00008	600000	3618	0.6%
cmp00009	347136	1815	0.5%
cmp00010	197250	513	0.3%
cmp00011	440238	2415	0.5%
cmp00012	369456	1294	0.4%
cmp00013	350889	1184	0.3%
cmp00014	269348	512	0.2%
cmp00015	292120	1018	0.3%
cmp00016	504828	1557	0.3%
cmp00017	346986	1117	0.3%

Table 8. JPEG 2000 Standard Transcoding Decoded with $r=0.5$.
Percentage of image area at each pixel difference is recorded.
No entry is shown when no pixels are at that error level.

Name	PixDiff =1	PixDiff =2	PixDiff =3	PixDiff =4	PixDiff =5	PixDiff =6
cmp00001	37%	5.4%	0.37%	0.017%	0.0006%	
cmp00002	22%	0.9%	0.02%	<0.000%		
cmp00003	11%	<0.0%				
cmp00004	14%	0.1%	<0.00%			
cmp00005	12%	0.0%	<0.00%			
cmp00006	26%	1.9%	0.10%	0.005%	0.0002%	0.0002%
cmp00007	29%	0.7%	0.02%			
cmp00008	24%	2.9%	0.21%	0.009%	0.0005%	
cmp00009	36%	6.2%	0.62%	0.048%	0.0023%	
cmp00010	37%	3.3%	0.14%	0.004%		
cmp00011	30%	1.0%	0.02%			
cmp00012	32%	2.7%	0.14%	0.005%		
cmp00013	37%	4.6%	0.25%	0.006%		
cmp00014	37%	3.9%	0.13%	0.003%		
cmp00015	19%	0.1%				
cmp00016	22%	0.2%	<0.00%			
cmp00017	38%	2.8%	0.07%	0.001%		

In this test, WSQ files generated by Aware's `wsq_demo` were transcoded to JPEG 2000 and then decoded using both the $r=0.56$ and $r=0.5$ options. IQM (v5.6) was computed on the reconstructed imagery from WSQ, the two standard transcoding options, and imagery that had been sequentially compressed with WSQ and JPEG 2000 Part 1 (WSQ-J2). A plot showing these results appears in Figure 7. It is clear that as expected the standard transcoding with $r=0.56$ decode is essentially identical to the WSQ results. The standard transcoding with $r=0.5$ decode gives somewhat reduced IQM results that very closely match or in a few instances lie just above the sequential compression alternative.




Figure 7. IQM Comparison of Transcoding Options

Alternate Transcoding

If the transcoded data will always be completely decoded without truncation, there is an alternate method for incorporating the effects of the non-standard reconstruction factor. If the transcoder anticipates that the decoder will be using " $r=0.5$ ", then it can incorporate the extra 0.06 adjustment into the zero bin shrinkage NZ. Reducing NZ by 0.06 in every subband achieves this. Notice that this changes only the transmitted parameters NZ, but the rest of the transcoded data remains identical. Therefore, there is no change in the transcoded file size.

When this alternate transcoding is used, the decoder reconstruction (using $r=0.5$) gives results very similar to an $r=0.56$ reconstruction on the standard transcoded file, as shown in Table 9. That is, the maximum difference in each output pixel is one gray-level, and the number of such differences is very close to that shown in Table 7.

Table 9. JPEG 2000 Alternate Transcoding Decoded with $r=0.5$.
Number of pixels off ground truth by 1 gray-level.

Name	#Pix_Total	Alternate JPEG 2000 $r=0.5$	%PixError
cmp00001	356345	1094	0.3%
cmp00002	638976	2616	0.4%
cmp00003	638976	2156	0.3%
cmp00004	612880	702	0.1%
cmp00005	638976	1702	0.3%
cmp00006	638976	4821	0.8%
cmp00007	347710	1649	0.5%
cmp00008	600000	3626	0.6%
cmp00009	347136	1816	0.5%
cmp00010	197250	510	0.3%
cmp00011	440238	2413	0.5%
cmp00012	369456	1293	0.4%
cmp00013	350889	1178	0.3%
cmp00014	269348	510	0.2%
cmp00015	292120	1018	0.3%
cmp00016	504828	1554	0.3%
cmp00017	346986	1117	0.3%

Although this transcoding alternative may increase the probability that the decoded image is as accurate as possible, it cannot guarantee it since JPEG 2000 decoders are free to adjust r to any value.

Compressed File Size

Sizes of the WSQ and JPEG 2000 compressed files were compared to see if there was any consistent difference. There are obviously some expected differences in header information. The JPEG 2000 file does not contain parameter listings for the (9,7)-filter coefficients, parameters R and C, or the Huffman table listings. And the WSQ file does not contain information specifying the tree decomposition used. Moreover, since the parameter syntax is different, parameters that are included take up less room in the JPEG 2000 file. A review of the extent of the header information in both formats shows that the JPEG 2000 header for these transcoded files is consistently 374 bytes, while the WSQ header ranges in size from 717 to 862 bytes. Therefore, the maximum header difference is on the order of 400 bytes.

Differences in header size, however, were not adequate to describe the differences in compressed file size shown in Table 10. The remainder of the difference is due to the increased compression efficiency of JPEG 2000. In general, the results show that at least a 10 percent improvement in compression efficiency is expected, compared to the WSQ compression found in the NIST fingerprint compression test set.

Table 10. Comparison of WSQ and JPEG 2000 Compressed File Size

Name	#Pix_Total	WSQ size	JPEG 2000 size	Diff	%Improvement JPEG 2000
cmp00001	356345	28114	24894	3220	11%
cmp00002	638976	41953	36344	5609	13%
cmp00003	638976	35158	30043	5115	15%
cmp00004	612880	32937	27957	4980	15%
cmp00005	638976	40299	34002	6297	16%
cmp00006	638976	39919	34665	5254	13%
cmp00007	347710	30584	27322	3262	11%
cmp00008	600000	33464	28391	5073	15%
cmp00009	347136	24805	21549	3256	13%
cmp00010	197250	16664	14662	2002	12%
cmp00011	440238	36833	32635	4198	11%
cmp00012	369456	30226	26505	3721	12%
cmp00013	350889	25293	22055	3238	13%
cmp00014	269348	20716	18010	2706	13%
cmp00015	292120	23749	20484	3265	14%
cmp00016	504828	33535	29149	4386	13%
cmp00017	346986	27609	24090	3519	13%

3.5.3 Visual Performance

When differences are as small as one gray-level, they are not perceptible to the human viewer. Differences as large as four or more gray-levels do have the potential of being perceived, so a small and admittedly non-rigorous visual inspection was performed. Flickering, as well as side-by-side visual inspection, was performed at 1x, 2x, and 4x pixel replication zoom of screen resolution.

First the reconstructed images with some of the largest discrepancies in the transcoding tests were inspected (cmp00009, cmp00006, and cmp00001). In no cases could any differences be seen between the ground truth reconstruction and the JPEG 2000 $r=0.5$ reconstruction (the transcoding option with the largest amount of error).

Since all the FBI compliance testing images were scans from inked cards, a further visual inspection was performed using livescan fingerprint imagery. In particular, print X_A2 from the image set used in the study described in Section 2 was chosen for inspection. During previous testing, it was possible to see image changes on this print when it was recompressed with JPEG 2000 Part 1, particularly when flickering was used.

For this additional X_A2 examination, a comparison was made between the image reconstructed with the Aware wsq_demo program and the one generated after a standard transcode to JPEG 2000 using $r=0.5$ decode. In this case, the maximum pixel difference was 9 gray-levels, and nearly 59 percent of the image pixels were at least one gray-level different from the Aware code reconstruction.

Visual inspection of these two images showed no perceptible differences in a side-by-side examination. Flicker tests at 4x zoom, however, did make very subtle fluctuations in gray-level visible. This slight flicker was so hard to notice that for any visual purposes the images were identical. In contrast, a flicker of WSQ and a sequentially coded WSQ-JPEG 2000 Part 1 image showed significant movement at the higher zoom levels. The same examination was performed using D5_377, with similar results except that no flickering at all could be seen between WSQ and the $r=0.5$ transcode. So, although the IQM of the $r=0.5$ transcode is very similar to that of the sequentially coded image (see Figure 7), the $r=0.5$ transcoding result is much closer to the WSQ result visually.

These examinations showed that the $r=0.5$ transcoded images were almost always visually indistinguishable from the WSQ reconstructions. In the rare cases where changes could be observed, they were extremely small, very hard to see, and had no impact on edge location.

Given the visual results, it would be reasonable to speculate that machine identification and matching using transcoded imagery (even using standard transcoding with $r=0.5$ decode) would more closely approximate results achieved using WSQ imagery than would a sequentially coded WSQ-JPEG2000 Part 1 image. However, visual inspection does not indicate how these differences from the ground truth imagery would impact machine interpretation of fingerprints, such as in an automated fingerprint identification system. Testing of the impact on machine interpretation was beyond the scope of this project.

3.6 JPEG 2000 to WSQ

Only transcoding from WSQ to JPEG 2000 has been implemented and tested. However, it is possible to make a few statements and inferences about transcoding in the opposite direction.

First, it must be recognized that although any FBI fingerprint WSQ file can be transcoded to JPEG 2000 Part 2, the converse is not true since there are a variety of options within JPEG 2000 that have no correlate in WSQ. However, within certain strict guidelines of quantization and wavelet filter and decomposition, it will be possible to transcode into WSQ. In particular, it is possible to take files that have been transcoded from WSQ to JPEG 2000 and turn them back into WSQ files.

When transcoding from JPEG 2000 to WSQ, the parameter precision does not decrease, so errors due to that aspect should not occur. However, there may still be small changes due to base 2 versus base 10 representation. The largest changes in the JPEG 2000 results were due to the uncertainty of the decoder reconstruction factor r . When transcoding into WSQ this uncertainty is removed. When performing the full loop, WSQ to JPEG2000 to WSQ, errors due to lost precision at the JPEG 2000 stage will still remain in the final WSQ file, but decoder uncertainty will not be a problem.

Also, since JPEG 2000 does not store the scaling value R explicitly, but rather incorporates it into the step sizes, the transcoded WSQ file will have $R=1$. This does not change the resultant image, but instead causes slight differences in the WSQ compressed representations in a full transcoding loop.

3.7 Conclusions

We have demonstrated that WSQ compressed files can be transcoded to JPEG 2000 Part 2 format with very small changes in reconstructed image pixel values. Although the reconstructed transcoding has very little, if any, visual change from the original reconstruction, it is outside the FBI-certification decoder requirements even when the FBI-specific decoder reconstruction factor $r=0.56$ is used.

Many of the small changes between a transcoded reconstruction and a WSQ reconstruction are inherent in the decreased precision of the JPEG 2000 step size representation and cannot be easily eliminated when transcoding from WSQ to JPEG 2000. However, larger changes were due to the imprecision of JPEG 2000 decoder. This issue can be addressed by setting guidelines for transcoding: specifying whether standard or alternative transcoding should be used and assigning a tag to identify the image type.

Finally, the JPEG 2000 transcoded file size was at least 10% smaller than the WSQ compressed file size on the test set used here. It should be borne in mind that this improvement in compressed size is offset to some extent by a reduction in the reconstructed pixel accuracy when using JPEG 2000.

Section 4

Summary

This study of WSQ and JPEG 2000 has revealed similarities and differences between the algorithms. This section summarizes some of the most important observations as they relate to the concepts of potential migration and coexistence.

- 1) By itself, JPEG 2000 Part 1 does not quite match the quality of WSQ when used on fingerprints. This difference seems to be primarily a result of the decomposition tree used. The Mallat decomposition available in JPEG 2000 Part 1 generates a slightly blockier looking reconstruction.
- 2) The sequential application of WSQ and JPEG 2000 Part 1 also has a somewhat blocky nature.
- 3) WSQ can be transcoded to JPEG 2000 Part 2 with no loss in visual quality for inked prints. For livescan impressions, in some cases the uncertainty in the JPEG 2000 Part 2 reconstruction just reached the threshold of perceptibility.
- 4) There are demonstrable, though very small, changes caused by the WSQ to JPEG 2000 Part 2 transcoding. There are enough differences to take the results outside the FBI certification guidelines. How these differences might affect automated fingerprint identification systems is unknown.
- 5) Transcoded JPEG 2000 Part 2 files are at least 10% smaller than the equivalent WSQ file.
- 6) A JPEG 2000 Part 2 encoder customized to choose step sizes in a manner similar to WSQ would not have the problems caused by reduced precision during transcoding. It is suspected that files produced by such an encoder would have quality very similar to WSQ. However, the imprecision of the reconstruction will still remain. This concept was not tested in this study.

Although JPEG 2000 Part 1 was not quite a visual match to WSQ, the quality may in fact be adequate for certain purposes since the visual appearance of ridge endings and bifurcations was not greatly altered. Since JPEG 2000 Part 1 software and hardware will be readily available in commercial products, decoding a WSQ file and recoding as JPEG 2000 Part 1 may be a convenient and adequate solution when a JPEG 2000 file is desired.

The transcoded JPEG 2000 Part 2 results are so similar to the original WSQ that they have potential to be a very good alternative when needing to communicate with a non-WSQ capable user. However, this presumes the source has access to a transcoder, and the destination has access to a JPEG 2000 Part 2 decoder that includes the variable offset, generalized scalar quantizer, and generalized decomposition capabilities.

JPEG 2000 Part 2 by itself may be capable of matching the quality of WSQ for fingerprint compression at a somewhat smaller file size. This naturally leads to questions about migration of the FBI fingerprint compression standard. However, this small improvement in file size must be weighed against other disadvantages of changing an accepted standard that is already in wide use.

Bibliography

1. R. Forkert and M. Lepley, Advanced Wavelet Image Compression, http://www.mitre.org/centers/cafc3/imagery_systems/imagelab/awic.html
2. Criminal Justice Information Services (CJIS) WSQ Gray-scale Fingerprint Image Compression Specification, Federal Bureau of Investigation, Document No. IAFIS-IC-0110 (V3), 19 Dec 1997.
3. ISO/IEC FCD15444-1, JPEG 2000 Part 1 Final Committee Draft Version 1.0, March 2000, <http://www.jpeg.org/FCD15444-1.htm>
4. JPEG 2000 Links, <http://www.jpeg.org/JPEG2000.htm>, many useful links to JPEG2000 documents and software
5. Image Quality Computation, <http://bmrc.berkeley.edu/courseware/cs294/fall97/assignment/psnr.html>
6. N. Nill, Image Quality Evaluation, <http://www.mitre.org/technology/mtf/>
7. ISO/IEC FCD15444-2, JPEG 2000 Part 2 Final Committee Draft, December 2000. Available to JPEG members via restricted access at <http://www.jpeg.org/private/register/wg1n2000.zip>
Anticipated public access <http://www.jpeg.org/FCD15444-2.htm>.
8. S Saha, Image Compression - from DCT to Wavelets : A Review, March 2000, <http://www.acm.org/crossroads/xrds6-3/sahaimgcoding.html>
9. S. Santini, Wavelet Tutorial Briefing, <http://www-cse.ucsd.edu/users/ssantini/cse228f/wavelet/wavelet.htm>
10. MathSoft Wavelet Resources, <http://www.mathsoft.com/wavelets.html> Links to many wavelet papers

Appendix A

Test Fingerprint Imagery

(see Table 2 for acquisition information)

card scan 1000 ppi



D1_125



M1_15



D1_377



M1_92

card scan 500 ppi



D5_125



D5_235



D5_377



D5_582

Appendix B

Image Quality Metric Data

Image	orig-contrast	orig-iqm	WSQ-psnr	WSQ-iqm	J2-psnr	J2-iqm
D1_125	0.27052	9956.983	30.646	9792.389	31.5726	8874.378
M1_15	0.21452	5287.012	31.3672	5225.952	32.7729	4876.469
D1_377	0.60525	50753.220	28.6955	51186.3	29.6923	46858.92
M1_92	0.71759	29882.040	29.5796	29395.84	30.3492	26852.46
D5_125	0.09664	432.857	35.2741	422.0836	36.3534	385.3156
D5_235	0.24026	1226.220	30.8157	1184.539	31.7777	1007.634
D5_377	0.63141	20915.980	25.8782	20642.13	26.837	18189.9
D5_582	0.78125	13724.940	27.3519	13453.63	28.3088	12351.06
H_WJ	0.1946	2769.340	27.0282	2519.191	28.2083	2323.513
H_KO	0.34376	7862.174	23.0411	7299.209	23.7413	6321.077
X_H2	0.61883	28505.830	22.7748	27348.15	23.5929	24525.99
X_A2	0.76298	36105.140	22.09	34117.83	22.8811	30406.61

Image	WSQ-WSQ-psnr	WSQ-WSQ-iqm	WSQ-J2-psnr	WSQ-J2-iqm	J2-WSQ-psnr	J2-WSQ-iqm
D1_125	30.5881	9497.121	29.6737	9449.692	30.1756	9146.205
M1_15	31.1277	5463.279	30.6419	5123.843	30.5377	5108.321
D1_377	28.4313	53262.95	27.7281	49566.44	28.4038	48198.06
M1_92	29.5872	28640.97	28.6418	28671.34	29.1551	28150.22
D5_125	35.0229	422.4271	34.3913	415.5557	34.6158	400.8916
D5_235	30.8243	1156.347	29.9911	1133.634	30.5565	1055.042
D5_377	25.8326	20758.19	25.0199	19752.96	25.7066	18920.23
D5_582	27.3765	12746.4	26.4075	13063.38	26.8654	12723.67
H_WJ	26.5624	2358.519	26.1857	2315.903	26.3699	2344.82
H_KO	22.8276	6904.231	21.9616	6585.656	22.6064	6464.717
X_H2	22.7146	25730.35	21.729	25753.78	22.4936	24894.51
X_A2	21.9601	34292.68	20.9565	31460.12	21.4869	31043.28

Image	J2-J2-psnr	J2-J2-iqm
D1_125	31.5582	8873.132
M1_15	32.7652	4880.413
D1_377	29.6921	46859.22
M1_92	30.3435	26877.17
D5_125	36.3505	385.3156
D5_235	31.7711	1006.623
D5_377	26.8353	18190.06
D5_582	28.2942	12326.98
H_WJ	27.9604	2351.672
H_KO	23.6248	6397.404
X_H2	23.5648	24521.02
X_A2	22.822	30320.82

	IQM Version 5.6			
Image	wsq	Standard Transcode $r=0.56$	Standard Transcode $r=0.5$	Wsq-JPEG2000 Part 1
D1_125	9792.40	9791.80	9392.10	9449.70
M1_15	5225.90	5225.70	4995.90	5123.80
D1_377	51186.00	51177.00	49271.00	49566.00
M1_92	29396.00	29393.00	28327.00	28671.00
D5_125	422.08	422.09	402.33	415.56
D5_235	1184.50	1184.50	1116.30	1133.60
D5_377	20642.00	20643.00	19698.00	19753.00
D5_582	13454.00	13454.00	12957.00	13063.00
H_WJ	2519.20	2519.10	2398.50	2315.90
H_KO	7299.20	7299.20	6908.80	6585.70
X_H2	27348.00	27349.00	26182.00	25754.00
X_A2	34118.00	34117.00	32509.00	31460.00

Appendix C

NIST Code Modification

Two technical bugs were found in the NIST code (wsq_v3_1) and corrected prior to the testing presented in this document. For the benefit of anyone trying to duplicate the results presented here or needing to use this version of the NIST code in future, an overview of these fixes is presented here. In addition, the corrected source files have been supplied to NIST. The next NIST WSQ release will fix these problems.

One of the bugs was quite minor and easy to fix. In the Huffman decoder, runlengths exactly equal to 100 were not decoded. This was easily fixed by changing the strict inequality check for runlengths (<100) to an inclusive inequality check (≤ 100) in the source code file huff.c

The next bug was much more extensive, but only impacted decoding for certain image sizes. In general terms the code that calculated subband sizes and computed wavelet filtering did not always allocate the appropriate number of elements to the lowpass and highpass filtered data.

More specifically, it was not recognized that the lowpass filtered data should always be at least as large as the highpass filtered data.

Fixing this second bug required that changes be made in the calculation of the wavelet and quantization tree structures and the wavelet filtering routines. The new functions for building wavelet and quantization structures include input flags indicating position reversal and these flags impact the subband length calculations. Also the wavelet filtering routines were modified so that filtering at the boundaries was independent of any inverted data positioning. Since the changes are somewhat complex, it is best to get an updated copy of the source code (tree.c and wsq_utils.c).

Although the appropriate testing has not been performed, initial runs seem to indicate that there is a good chance that this modified version of wsq_v3_1 could meet FBI certification tests.

Glossary

(9,7)	Indication of wavelet filter (lowpass, highpass) lengths
AFIS	Automated fingerprint identification system
bit-plane	A two-dimensional array of bits of the same magnitude from all coefficients
bitrate	Average bits per pixel in the compressed file
bitstream	The actual sequence of bits resulting from the coding of a sequence of symbols prior to being formatted with a particular syntax
coding	Lossless remapping of data that typically occupies less space: types include Huffman, arithmetic, DPCM, runlength, etc.
compression	Series of techniques that in combination generally makes image storage smaller. Typically includes transform+quantizer+encoder+syntax.
compression ratio	Ratio of original file size to compressed file size
decomposition	Exact structure of wavelet application. Types include Mallat and packet.
decompression	Series of techniques that in combination produce a two-dimensional array of values from compressed image data
FBI	Federal Bureau of Investigation
ground truth	Best possible results/data. Used when testing algorithms.
ISO	International Organization for Standardization
JPEG	Joint Photographic Experts Group
LL band	Lowpass band (subsampling version of the original data)
LH, HL, HH bands	Highpass bands containing edge information at different resolutions with the band name indicating general edge orientation
lossless	The effect of an overall compression/decompression process in which the output of the decompression is identical to the input of the encoding process
lossy	The effect of an overall compression/decompression process in which the output of the decompression is not identical to the input of the encoding process
marker (segment)	A two-byte code labeling the type of information and following data
ppi	Pixels per inch

progression order	The order of transform coefficient information in the bitstream
quantization	Operation to reduce precision of individual coefficients
resolution level	Subbands that when reconstructed produce an image of a particular resolution relative to the original
scan order	Order in which coefficients are processed during coding
step size	Parameter specifying the degree of quantization
subband	A group of transform coefficients resulting from the same sequence of lowpass and highpass operations, both vertically and horizontally
syntax	File format structures, marker design, tags, etc., that specify exactly how compressed image data is stored in a file
transcode	Partially decode one format and recode into another format
wavelet filter	A filter pair that used to generate a wavelet transform
wavelet transform	Mathematical operation that decomposes an image into multiresolution coefficients consisting of lowpass (reduced resolution) and highpass (scaled edge information) parts
WSQ	Wavelet Scalar Quantization
zero bin	Quantized data mapped to zero